

11/09/00  
JC600 U.S. PTO  
Frederick P. Fish  
1855-1930  
W.K. Richardson  
1859-1951  
BOSTON  
DALLAS  
DELAWARE  
NEW YORK  
SAN DIEGO  
SILICON VALLEY  
TWIN CITIES  
WASHINGTON, DC

11-14-00

A

# FISH & RICHARDSON P.C.

VIA EXPRESS MAIL

November 9, 2000

225 Franklin Street  
Boston, Massachusetts  
02110-2804

Telephone  
617 542-5070

Facsimile  
617 542-8906

Web Site  
www.fr.com

Attorney Docket No.: US-4190-CIP/10454-014002

**Box Patent Application**  
Commissioner for Patents  
Washington, DC 20231

JC600 U.S. PTO  
09/711323  
11/09/00

Presented for filing is a new continuation-in-part patent application of:

**Applicant:** ALFONSO VALDES, KETIH SKINNER AND PHILLIP ANDREW PORRAS

**Title:** SENSOR AND ALERT CORRELATION IN INTRUSION DETECTION SYSTEMS

Enclosed are the following papers, including those required to receive a filing date under 37 CFR §1.53(b):

	Pages
Specification	10
Claims	3
Abstract	1
Declaration	2 [Unsigned]
Drawing(s)	4

Enclosures:

- Postcard.
- Appendix (19 pages)

This application is a continuation-in-part (and claims the benefit of priority under 35 USC 120) of U.S. application serial no. 09/653,066, filed September 1, 2000. The disclosure of the prior application is considered part of (and is incorporated by reference in) the disclosure of this application.

## CERTIFICATE OF MAILING BY EXPRESS MAIL

Express Mail Label No. EL 624 275 981 US

I hereby certify under 37 CFR §1.10 that this correspondence is being deposited with the United States Postal Service as Express Mail Post Office to Addressee with sufficient postage on the date indicated below and is addressed to the Commissioner for Patents, Washington, D.C. 20231.

Date of Deposit

November 9, 2000

Signature

Samantha Bell  
Samantha Bell

Typed or Printed Name of Person Signing Certificate

FISH & RICHARDSON P.C.

Commissioner for Patents

November 9, 2000

Page 2

This application is entitled to small entity status. A small entity statement will be filed at a later date.

Basic filing fee	\$355
Total claims in excess of 20 times \$9	\$0
Independent claims in excess of 3 times \$40	\$120
Fee for multiple dependent claims	\$0
Total filing fee:	\$475

A check for the filing fee is enclosed. Please apply any other required fees or any credits to deposit account 06-1050, referencing the attorney docket number shown above.

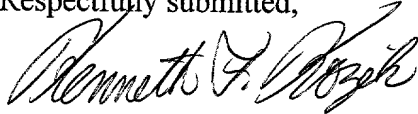
If this application is found to be incomplete, or if a telephone conference would otherwise be helpful, please call the undersigned at (617) 542-5070.

Kindly acknowledge receipt of this application by returning the enclosed postcard.

Please send all correspondence to:

DAVID L. FEIGENBAUM  
Fish & Richardson P.C.  
225 Franklin Street  
Boston, MA 02110-2804

Respectfully submitted,



Kenneth F. Kozik  
Reg. No. 36,572

Enclosures  
KFK/dmm  
20154936.doc

# **APPLICATION FOR UNITED STATES PATENT**

**in the name of**

**Alfonso Valdes, Ketih Skinner and  
Phillip Andrew Porras**

**of**

**SRI International, Inc.**

**for**

**Sensor and Alert Correlation in Intrusion Detection  
Systems**

**Kenneth F. Kozik**  
Fish & Richardson P.C.  
225 Franklin Street  
Boston, MA 02110-2804  
Tel.: (617) 542-5070  
Fax: (617) 542-8906

**ATTORNEY DOCKET:**  
**BOS-10454-KFK110900-1**

**DATE OF DEPOSIT:**

November 9, 2000

**EXPRESS MAIL NO.:**

EL 624 275 981 US

## Sensor and Alert Correlation in Intrusion Detection Systems

### Reference to Government Funding

This invention was made with Government support under contract number  
5 F30602-99-C-0149 awarded by the Air Force Research Laboratory. The Government has  
certain rights in this invention.

### Related Application

This is a continuation-in-part of U.S. patent application serial number 09/653,066,  
filed September 1, 2000 and entitled "Methods for Detecting and Diagnosing  
10 Abnormalities Using Real-Time Bayes Networks," which is incorporated herein by  
reference.

### Reference to Related Documents

This invention incorporates by reference a paper entitled "Probabilistic  
Approaches to Alert Management," by A. Valdes, attached as an appendix.  
15

### Background

The invention relates generally to intrusion detection, and more specifically to  
sensor and alert correlation in intrusion detection systems.

There are three main types of intrusion detection systems currently used to detect  
20 hacker attacks on computer networks: signature analysis systems, statistical analysis  
systems, and systems based on probabilistic reasoning. Signature analysis systems  
compare current data traffic patterns with stored traffic patterns representing the signature  
or profile of various types of hacker attacks. These systems generate an alert if the pattern  
of traffic received by the network matches one of the stored attack patterns.

25 Statistical analysis systems compare current data traffic patterns with statistical  
profiles of previous traffic patterns. These systems generate an alert if a current traffic  
pattern is significantly different from a stored profile of "normal" traffic. Examples of  
both statistical- and signature-based intrusion detection systems are described in Porras,  
et al., "Live Traffic Analysis of TCP/IP Gateways," Internet Society's Networks and  
30 Distributed Systems Society Symposium, March 1998.

Examples of intrusion detection systems based on probabilistic reasoning are described in U.S. patent application serial number 09/653,066 entitled "Methods for Detecting and Diagnosing Abnormalities Using Real-Time Bayes Networks." In one such system, a Bayes network is established that includes models (called "hypotheses")

5 that represent both normal traffic and attack traffic received by a computer network. Traffic actually received by the network is examined in real-time to identify its relevant characteristics or features, such as volume of data transfer, number of erroneous connection requests, nature of erroneous connection requests, ports to which connections are attempted, etc. Information about these relevant features is then provided to the

10 Bayes network, which calculates a system belief (a probability) that the current network traffic is either normal traffic or attack traffic.

A typical intrusion detection system may include one or more sensors that monitor network traffic in the manner discussed above, and one or more other sensors that monitor network resources. A system operator or network administrator (usually a

15 person) reviews all of the alerts generated by the system.

A major problem with existing intrusion detection systems is that they often provide misleading, incomplete, or low-quality information to the system operator; this may make it impossible to take proper steps to protect the network. For example, during a large-scale hacker attack each of the system's sensors may generate hundreds of alerts.

20 Although each alert may be accurate, the sheer number of alerts could easily overwhelm the system operator. Moreover, false alarms can be triggered by normal traffic directed towards a malfunctioning network resource, such as a server. These false alarms could distract the system operator from alerts triggered by actual hacker attacks. Finally, most intrusion detection systems are unable to detect low-level attacks such as port sweeps, in

25 which hackers slowly "probe" a network to discover its structure and weaknesses.

### Summary

This invention uses probabilistic correlation techniques to increase sensitivity, reduce false alarms, and improve alert report quality in intrusion detection systems. In

30 one preferred embodiment, an intrusion detection system includes at least two sensors to monitor different aspects of a computer network, such as a sensor that monitors network

traffic and a sensor that discovers and monitors available network resources. The sensors are correlated in that the belief state of one sensor is used to update or modify the belief state of another sensor. For example, a network resource sensor may detect that a server is malfunctioning. This information is used to modify the belief state of one or more  
5 network traffic sensors so that normal network traffic directed towards the malfunctioning server does not trigger a false alarm.

In another example, a network resource sensor transmits its knowledge of network structure to a network traffic sensor. This information is used to modify the belief state of the network traffic sensor so that an attempt to communicate with a non-existent  
10 resource appears to be suspicious. By allowing different sensors to share information, the system's sensitivity to low-level attacks can be greatly increased, while at the same time greatly reducing the number of false alarms.

In another embodiment of this invention, probabilistic correlation techniques are used to organize alerts generated by different types of sensors. By comparing features of  
15 each new alert with features of previous alerts, and adjusting the comparison by an expectation that certain feature values will or will not match, the alerts can be grouped in an intelligent manner. The system operator may then be presented with a few groups of related alerts, rather than dozens of individual alerts from multiple sensors.

## Brief description of the drawings

**Figure 1** shows a preferred intrusion detection system with coupled sensors.

**Figure 2** is a flowchart showing a preferred method for coupling sensors.

**Figure 3** shows a preferred intrusion detection system with coupled sensors and an alert correlation device.

**Figure 4** is a flowchart showing a preferred method for grouping alerts into classes of related alerts.

## Detailed Description

### I. Introduction

In preferred embodiments, intrusion detection systems for computer networks include sensors that monitor both network traffic and network resources. Correlation techniques are used to increase system sensitivity, reduce false alarms, and organize alerts into related groups. The sensor and alert correlation techniques described below may be used in intrusion detection systems that include any type or mix of sensors. However, probabilistic sensors similar to those described in U.S. Patent Application Serial Number 09/653,066 may provide advantages over other types of sensors.

Bayes networks are preferably used to perform the probabilistic reasoning functions described below. As is more fully described in U.S. Patent Application Serial Number 09/653,066, Bayes networks include models (called “hypotheses”) that represent some condition or state; probabilistic reasoning methods are used to determine a belief that a current observation corresponds to one of the stored hypotheses. Before a Bayes network performs an analysis, an initial belief (called a “prior probability”) regarding the various hypotheses is established. During the analysis, a new belief based on actual observations is established, and the prior probability is adjusted accordingly. In the context of the embodiments described here, a relevant Bayes hypothesis might be “there is a stealthy portsweep attack against the computer network.” Belief in the portsweep attack hypothesis would be strengthened if attempts to connect to several invalid or nonexistent ports are observed.

### II. Intrusion Detection with Correlated Sensors

As was discussed above, intrusion detection systems typically have several independent sensors that gather different types of information. In preferred embodiments of this invention, the information gathered by the various sensors may be shared to improve overall performance of the intrusion detection system.

In a preferred intrusion detection system 100 shown in Figure 1, sensors are coupled so that the belief state of one sensor (such as network resource sensor 103) affects the belief state of another sensor (such as network traffic sensor 101). Each

sensor may then transmit alerts to a system operator or network administrator 105. These sensors need not be probabilistic sensors; for example, a sensor that keeps track of the number of servers in a network has wouldn't need to generate a probabilistic output.

Figure 2 illustrates a preferred method by which sensors are coupled or correlated.

5 At step 201, a first (preferably probabilistic) sensor receives all or part of the belief state of a second (not necessarily probabilistic) sensor. The belief state of the second sensor may indicate an apparent normal, degraded, or compromised state of a monitored system resource, the existence or validity of supported services, or any other relevant belief state held by a sensor in an intrusion detection system. At step 203, a prior belief state of the  
10 first sensor is adjusted, the adjustment based at least in part on the second sensor's belief state. For example, the prior belief state of the first sensor or may be adjusted so that an erroneous transaction with a damaged or compromised network resource does not generate an alert. In another example, the prior belief state of the first sensor may be adjusted so that an attempted communication with a system server or resource that does  
15 not exist appears to be suspicious. By allowing different sensors to share information, the system's sensitivity to low-level attacks can be greatly increased, while at the same time greatly reducing the number of false alarms.

### III. Alert Correlation

20 In another embodiment of this invention, probabilistic correlation techniques are used to organize alerts into classes or groups of related alerts. By comparing features of each new alert with features of previous alerts, and adjusting the comparison by an expectation that certain feature values will or will not match, the alerts can be grouped in an intelligent manner. Alerts may be grouped or organized in several different ways,  
25 depending on both the type of an attack and the structures of the intrusion detection system and the monitored network. For example, all alerts related to a specific attack (such as a denial of service attack) may be grouped together. Or, all alerts related to the various stages of a staged attack may be grouped together to allow the system operator to view the progression of a staged attack.



### A. Alert Correlation Devices

Figure 3 shows a preferred intrusion detection system 300 that includes a network traffic sensor 301 and network resource sensor 303. Intrusion detection system 300 may include any number of sensors, and the belief states of the sensors may be correlated in the manner discussed above. Alerts generated by sensors 301 and 303 are provided to an alert correlation device 305, which uses probabilistic reasoning (preferably Bayesian) techniques to organize alerts into classes of related alerts. These alert classes are then provided to a system operator or network administrator 307.

### B. Measurement of Similarity

To group alerts in an intelligent manner, it is necessary to define and determine the degree of similarity between a new alert and an existing alert or alert class. This measurement of similarity is preferably calculated by comparing similarity among shared features (also called “measures”) of the alert and alert class. Examples of features to consider when assessing similarity between a new alert and an existing alert class include source IP address, destination IP address and port, type of alert, type of attack, etc. The nature of an alert may change the expectation of which features should be similar. It is therefore important to define a measurement of similarity to take into account which features are candidates for matching.

In the following discussion, X is defined to be the value of a feature or measure in a new alert, and Y is the value of that feature or measure in an existing alert class to which it is being compared. For features that have a single value, the features of a new alert and an existing alert class are defined to be similar if their feature values are equal. That is:

$$SIM(X, Y) = \begin{cases} 1.0, & X = Y \\ 0, & otherwise \end{cases}$$

The more general case is one in which a feature of an alert includes a list of observed values, and a “hit count” of the number of times each value has been observed. For reasons of normalization, the hit count may be converted to a probability. In this sense, X and Y are lists of feature and probability values, possibly of different lengths. A probability vector describes a pattern over observed categories, where:

$p_X(C)$  = probability of category  $C$  in list  $X$ ,

$p_Y(C)$  = probability of category  $C$  in list  $Y$ ,

$\mathbf{P}_X$  = probability vector over categories observed for  $X$ ,

$\mathbf{P}_Y$  = probability vector over categories observed for  $Y$ .

The notation  $C \in X$  is used to denote that category  $C$  occurs in list  $X$ . The similarity of the two lists is given by:

$$Sim(X, Y) = \frac{\left[ \sum_{C \in X \text{ AND } C \in Y} P_X(C) \times P_Y(C) \right]^2}{(\mathbf{P}_X \bullet \mathbf{P}_X)(\mathbf{P}_Y \bullet \mathbf{P}_Y)}$$

5 If the two lists are the same length, then this measure gives the square of the cosine between the two. This has an intuitive geometric property that is not shared by, say, the dot product as is commonly used in the pattern matching literature.

If the patterns  $(1, 0)$ ,  $(0, 1)$ , and  $(0.5, 0.5)$  are over the same two features, then:

$$Sim(\{1, 0\}, \{0, 1\}) = 0$$

$$Sim(\{0.5, 0.5\}, \{0, 1\}) = Sim(\{0.5, 0.5\}, \{1, 0\}) = 0.5$$

10 In other words, the first two patterns are orthogonal, and the third is halfway in between.

### C. Expectation of Similarity

As was discussed above, the nature of an alert may change an expectation of which features of a new alert and an existing alert class should be similar. For example, a syn flood is a type of attack in which the source IP address is typically forged. In this case, similarity in the source IP address would not be considered when assessing the overall similarity between a new alert triggered by a syn flood attack and an existing alert class.

15 In a preferred embodiment, the expectation of similarity is a feature-specific number between 0 and 1, with 1 indicating a strong expectation of a match, and 0 indicating that two otherwise similar alerts are likely not similar with respect to a specific feature.

25 An alert state is preferably represented as a probability vector over likely alert states. For each candidate alert state, each feature has a vector of the same length whose elements are the similarity expectation values (that is, numbers between 0 and 1) given

the state. These expectation values may initially be assigned a value of about 0.6 to indicate a medium expectation of similarity, except where it is known that the expectation of similarity is either lower or higher. For example, in an access from a compromised host or a denial of service (DOS) attack, the attacker often spoofs (makes up) the source IP address. Therefore, the expectation values may initially be assigned a value of about 0.1, indicating a low expectation that the source IP address will match that of the attacker.

Based on the alert state, the similarity expectation may be dynamically generated as the weighted sum of the elements of an expectation table, with the weights from the (evolving) alert state distribution. This is itself an expectation in the statistical sense, conditioned on the belief over the present alert state. The similarity expectation is therefore general enough to cover the situation of an unambiguous alert or “call” from a signature-based sensor, in which case the distribution over states is 1.0 at the state corresponding to the alert or call and 0 elsewhere. For example, algebraically for a feature or measure  $J$ :

$$E_J = \sum_{i \in \{alert\_states\}} BEL(i) E_J(i)$$

$E_J$  = Expected similarity for measure  $J$ , given present state.

$BEL(i)$  = Belief that the attack state is presently  $i$ .

$E_J(i)$  = Element  $i$  of the lookup table of expected similarity for measure  $J$  given state  $i$ .

Both the new alert and the alert class to which it is being compared each compute the similarity expectation for each feature. Feature similarity and similarity expectation are then combined to form a single value of alert similarity. This is done by combining feature similarities and normalizing by similarity expectation, over the set of common features. For example:

$$SIM(X, Y) = \frac{\sum_J E_J^X E_J^Y SIM(X_J, Y_J)}{\sum_J E_J^X E_J^Y}$$

$X$  = New alert

$Y$  = Existing alert

$J$  indexes measures

$E_J^{X(Y)}$  = Similarity expectation for measure  $J$ , alert  $X(Y)$ .

A similar definition can be formed by taking products rather than sums above, and then taking the geometric mean. This has some advantages, but can be overly influenced by a single extremely small value.

Where appropriate, similarity expectation may also cover list containment and subnet similarity. For example, an intrusion detection system may generate two alerts, one from a network sensor and one from a host sensor. To decide if the target addresses of these two alerts are similar, the expectation of similarity might be that the target address in the alert generated by the host sensor would be contained in the list of target addresses in the network sensor's alert. In another example, an intermediate expectation of similarity may be generated if target addresses from two alerts do not match but appear to be from the same subnet.

#### **D. Transition Models**

When a new alert is encountered, all existing alert classes are preferably passed through their transition models to generate new prior belief states. Transition models attempt to assign a "time value" to alert confidence, as well as try to anticipate the next step in a staged attack. The "time value" is typically modeled as a multiplicative decay factor that reduces alert model confidence slowly over a period of days. The decay period varies by type of attack, and eventually decreases to some background level. This is used to downweight very old alert classes when there are multiple candidates to be compared with a new alert. Transition in time tends to decrease overall suspicion for a specific alert class, as well as spread suspicion over other hypotheses. That is, it tends to make the decision whether to group a new alert with an existing alert class less confident.

The ability to anticipate the next stage in a staged attack is also achieved by a transition in state. Internally, a transition matrix is maintained that expresses the probability of the next state in a staged attack, given the present state. Persistence (the fact that the next step may match the current) is explicitly captured. The transition gives a prior distribution over attack states that is assumed to hold immediately before the next alert is analyzed.

After the transition models generate new prior belief states for their respective alert classes, the similarity expectation for each feature in a new alert is updated. The similarity of the new alert to all existing alert classes is then computed. If none match

above some similarity threshold, the new alert is defined to be a new alert class.

Otherwise, it is assumed that the new alert is a continuation of the best-matching alert class. By including a list of contributing alerts in the alert class structure, full details of the new alert are retained.

## 5           **E. Alert Correlation Method**

Figure 4 is a flowchart showing how alerts may be aggregated into classes of related alerts in a preferred embodiment. As was discussed above, both alerts and alert classes each have one or more distinguishing features that can be assigned different values. First, a set of potentially similar features shared by a new alert and one or more  
 10 existing alert classes is identified (step 401). Next, an expectation of similarity between the features of the new alert and features of one or more existing alert classes is either generated or updated (step 403). After a comparison between the new alert and the existing alert class(es) is complete (step 405), the new alert is either associated with the existing alert class that it most closely matches (step 407), or new alert class is defined to  
 15 include the new alert (step 409).

Although preferred embodiments have been discussed above, the scope of this invention is defined by the following claims:

**What is claimed is:**

1. A method for correlating a first sensor to a second sensor in an intrusion detection system, the first and second sensors each maintaining belief over a number of possible states of the system, the method comprising the steps of:

- 5           (a) transmitting to the first sensor information about the second sensor's belief state; and
- (b) adjusting a prior belief state of the first sensor, the adjustment based at least in part on the second sensor's belief state.

2. The method of claim 1 wherein the first and second sensors are different types of  
10           sensors.

3. The method of claim 2 wherein the first sensor is a probabilistic sensor.

4. A method for reducing false alarms generated by an intrusion detection system when a monitored resource is degraded or compromised, the intrusion detection system having a first and second sensors each maintaining belief over a number of possible  
15           states of the system, the method comprising the steps of:

- (a) transmitting to the first sensor all or part of the belief of the second sensor regarding an apparent normal, degraded or compromised state of a monitored resource; and
- (b) adjusting a prior belief state of the first sensor so that an erroneous transaction  
20           with the degraded or compromised resource does not generate an alarm.

5. A method for enhancing the sensitivity of an intrusion detection system that monitors a plurality of computer system resources, the intrusion detection system having a first and second sensors each maintaining belief over a number of possible states of the system, the method comprising the steps of:

(a) transmitting to the first sensor all or part of the belief of the second sensor regarding the existence or validity of services supported on monitored computer system resources; and

(b) adjusting a prior belief state of the first sensor so that an attempted communication with a nonexistent system service or resource appears suspicious.

6. A method for organizing alerts into alert classes, both the alerts and alert classes having a plurality of features, each feature having one or more values, the method comprising the steps of:

(a) identifying a set of potentially similar features shared by a new alert and one or more existing alert classes;

(b) comparing the new alert to one or more existing alert classes;

(c) adjusting the comparison by an expectation that certain feature values will or will not match, and either:

(d1) associating the new alert with the existing alert class that the new alert most closely matches; or

(d2) defining a new alert class that is associated with the new alert.

7. A method for organizing alerts into alert classes, both the alerts and alert classes having a plurality of features, each feature having one or more values, the method comprising the steps of:

(a) receiving a new alert;

(b) identifying a set of potentially similar features shared by the new alert and one or more existing alert classes;

(c) updating a similarity expectation for one or more feature values;

(d) comparing the new alert with one or more alert classes, and either:

(e1) associating the new alert with the existing alert class that the new alert most closely matches; or

(e2) defining a new alert class that is associated with the new alert.

8. The method of claim 7 further comprising the step (a1) of passing each existing alert class through a transition model to generate a new prior belief state for each alert class.

9. A method for organizing alerts having a plurality of features, each feature having one or more values, the method comprising the steps of:

(a) generating a group of feature records for a new alert, each feature record including a list of observed values for its corresponding feature;

(b) identifying a set of potentially similar features shared by the new alert and one or more existing alert classes that are associated with previous alerts;

(c) comparing the new alert to one or more alert classes;

(d) adjusting the comparison by an expectation that certain feature values will or will not match, and either:

(e1) associating the new alert with the existing alert class that the new alert most closely matches; or

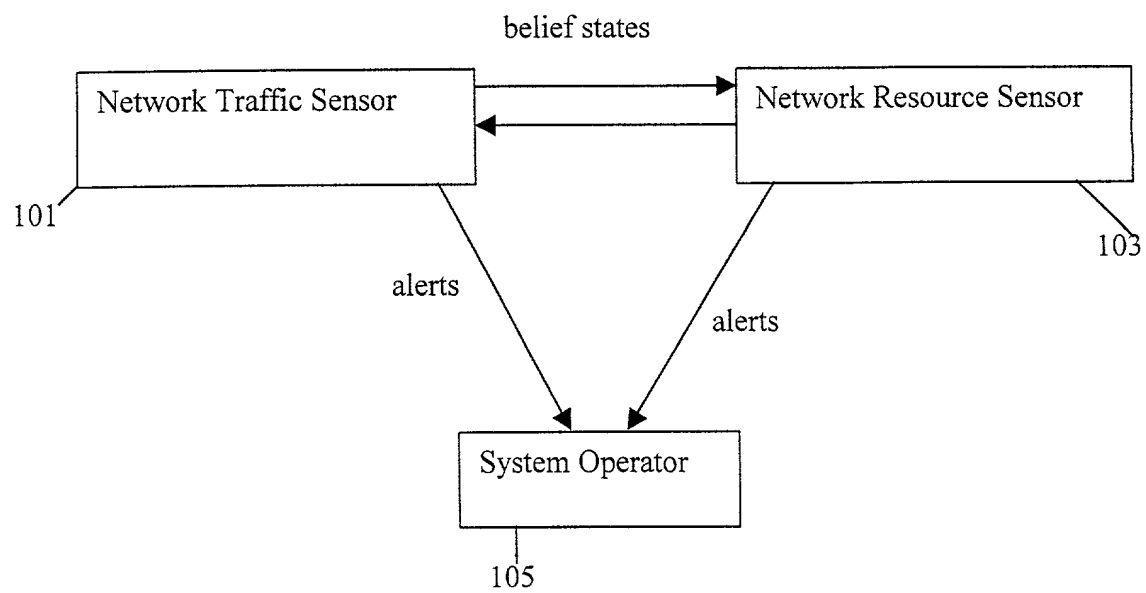
(e2) defining a new alert class that is associated with the new alert.



**ABSTRACT**

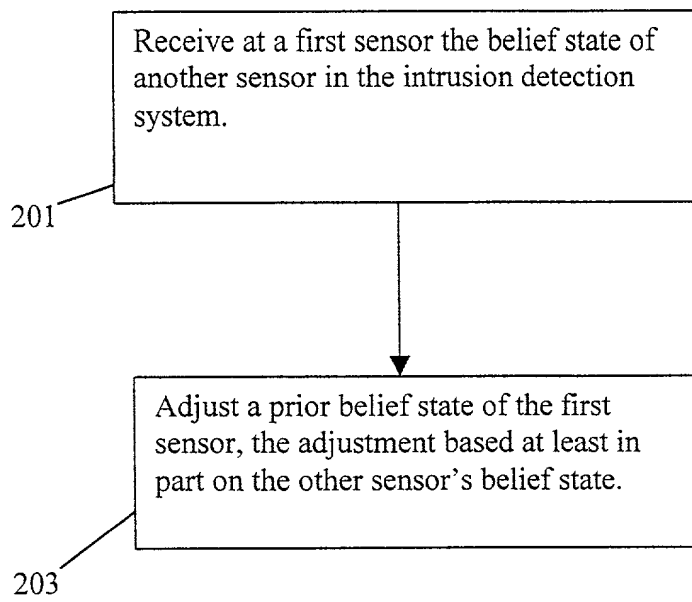
This invention uses probabilistic correlation techniques to increase sensitivity, reduce false alarms, and improve alert report quality in intrusion detection systems. In one preferred embodiment, an intrusion detection system includes at least two sensors to monitor different aspects of a computer network, such as a sensor that monitors network traffic and a sensor that discovers and monitors available network resources. The sensors are correlated in that the belief state of one sensor is used to update or modify the belief state of another sensor. In another embodiment of this invention, probabilistic correlation techniques are used to organize alerts generated by different sensors in an intrusion detection system. By comparing features of each new alert with features of previous alerts, and adjusting the comparison by an expectation that certain feature values will or will not match, the alerts can be grouped in an intelligent manner.

20169104.doc

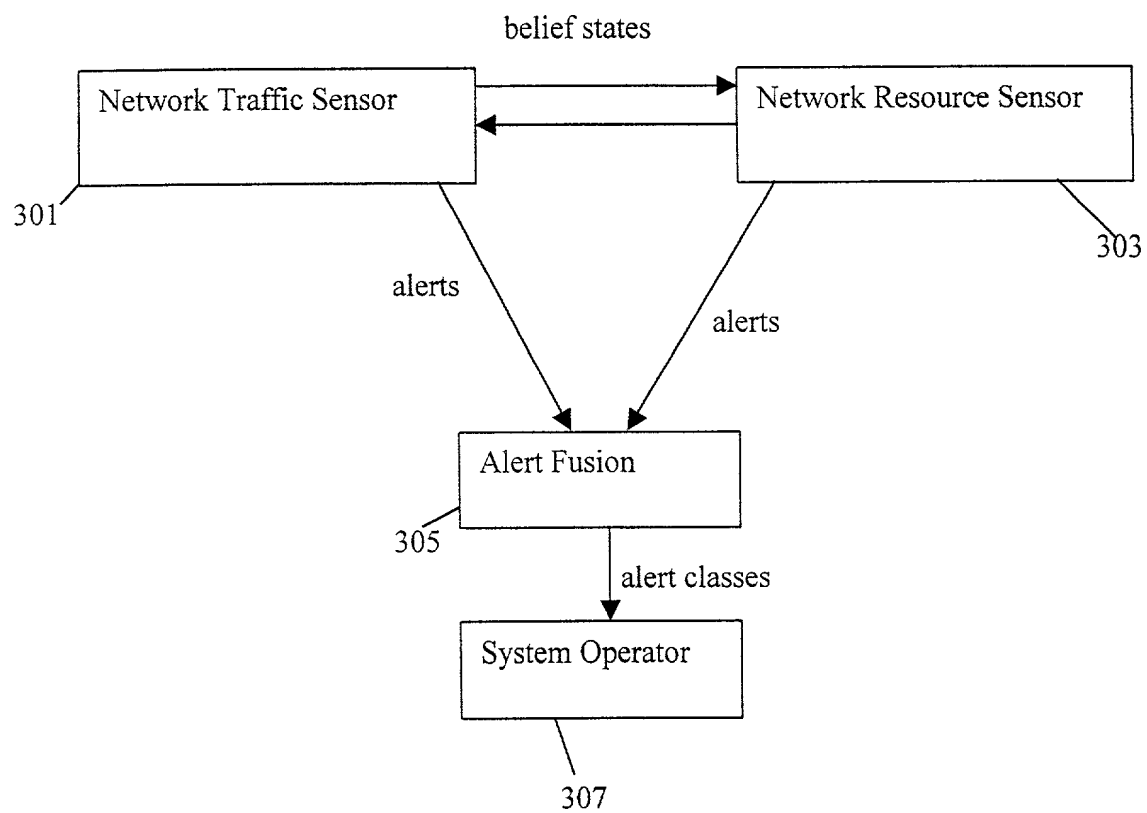


100

**FIGURE 1**



**FIGURE 2**



300

**FIGURE 3**

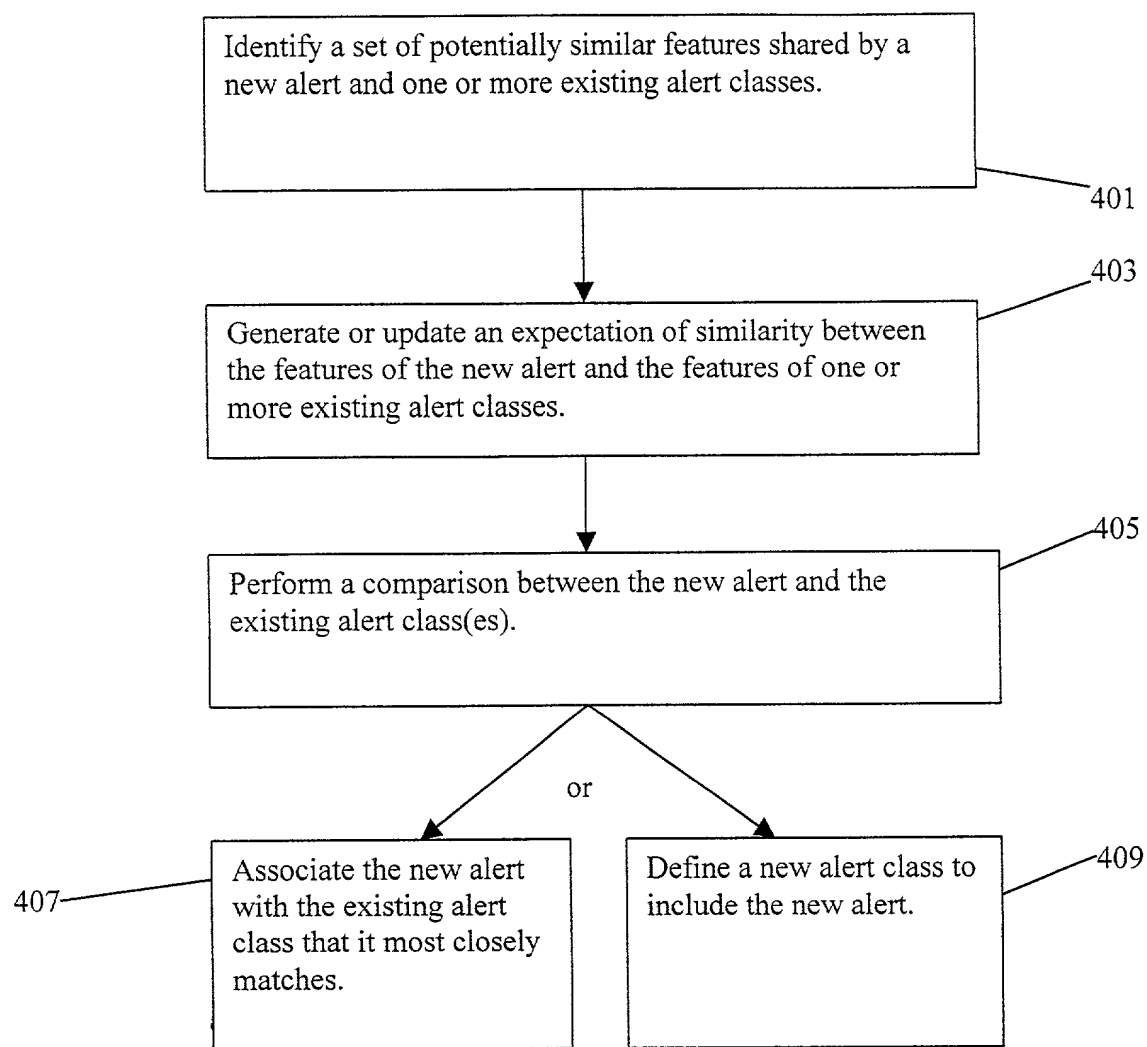


FIGURE 4

# COMBINED DECLARATION AND POWER OF ATTORNEY

As a below named inventor, I hereby declare that:

My residence, post office address and citizenship are as stated below next to my name.

I believe I am the original, first and sole inventor (if only one name is listed below) or an original, first and joint inventor (if plural names are listed below) of the subject matter which is claimed and for which a patent is sought on the invention entitled SENSOR AND ALERT CORRELATION IN INTRUSION DETECTION SYSTEMS, the specification of which:

- ☒ is attached hereto.  
☐ was filed on \_ as Application Serial No. \_ and was amended on \_\_\_\_\_.  
☐ was described and claimed in PCT International Application No. \_\_\_\_\_ filed on \_\_\_\_\_ and as amended under PCT Article 19 on \_\_\_\_\_.

I hereby state that I have reviewed and understand the contents of the above-identified specification, including the claims, as amended by any amendment referred to above.

I acknowledge the duty to disclose all information I know to be material to patentability in accordance with Title 37, Code of Federal Regulations, §1.56.

I hereby claim the benefit under Title 35, United States Code, §120 of any United States application(s) listed below and, insofar as the subject matter of each of the claims of this application is not disclosed in the prior United States application in the manner provided by the first paragraph of Title 35, United States Code, §112, I acknowledge the duty to disclose all information I know to be material to patentability as defined in Title 37, Code of Federal Regulations, §1.56(a) which became available between the filing date of the prior application and the national or PCT international filing date of this application:

U.S. Serial No.	Filing Date	Status
09/653,066	Sept. 1, 2000	Pending

I hereby appoint the following attorneys and/or agents to prosecute this application and to transact all business in the Patent and Trademark Office connected therewith:

David L. Feigenbaum, Reg. No. 30,378	Kenneth F. Kozik, Reg. No. 36,572
Steven Weiner, Reg. No. 38,360	Edward Davis, Reg. No. 35,112

Address all telephone calls to DAVID L. FEIGENBAUM at telephone number (617) 542-5070.

Address all correspondence to DAVID L. FEIGENBAUM at:

FISH & RICHARDSON P.C.  
 225 Franklin Street  
 Boston, MA 02110-2804

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application or any patents issued thereon.

## Combined Declaration and Power of Attorney

Page 2 of 2 Pages

Full Name of Inventor: Alfonso de Jesus Valdes

Inventor's Signature: \_\_\_\_\_ Date: \_\_\_\_\_  
Residence Address: 2223 Carmelita Drive, San Carlos, CA 94070  
Citizenship: USA  
Post Office Address: same as above

Full Name of Inventor: Keith Michael Skinner

Inventor's Signature: \_\_\_\_\_ Date: \_\_\_\_\_  
Residence Address: 912 Suntree Court, Sunnyvale, CA 94086  
Citizenship: USA  
Post Office Address: same as above

Full Name of Inventor: Phillip A. Porras

Inventor's Signature: \_\_\_\_\_ Date: \_\_\_\_\_  
Residence Address: 762 Ormonde Drive, Mountain View, CA 94043  
Citizenship: USA  
Post Office Address: same as above

20180187.doc

# Probabilistic Approaches to Alert Management

## Abstract

With the growing deployment of host and network intrusion detection systems, managing reports from these systems becomes critically important. Alert reports contain a variety of attributes describing, for example, the sensor's confidence in the attack, the nature of the attack, and the assets affected. Generally, not all attributes are specified on any given alert report, and management approaches should make optimal use of the attributes available. We present two components that employ probabilistic approaches to address two key problems in alert management: alert correlation and prioritization. For correlation, we extend ideas from multisensor data fusion. For shared attributes between alerts that are candidates for fusion, we consider appropriate similarity measures adjusted for situation-specific expectation of similarity. Prioritization is addressed in a Bayes model that produces an alert ranking in agreement with a human expert. By assuming that dependencies between attributes are limited to attribute groupings, the required model can be specified compactly. Our model comprehends differential weighting of attributes in an alert, incorporation of preference profiles, optimal use of observed attributes, update capability, and extensibility. We present the base model as well as a variant that can be interactively trained by a human expert. We also present results from an initial deployment of these components.

**Keywords:** Network security, distributed system security, alert management, alert prioritization, adaptive systems

JCS00 U.S. PTO  
09/711323  
11/09/00



## Introduction

In response to attacks and potential attacks against enterprise networks, administrators are increasingly deploying intrusion detection systems (IDSs). These systems monitor hosts, networks, critical files, and so forth, using a variety of signature and probabilistic techniques [1, 2]. The use of such systems has given rise to another difficulty, namely, the intelligent management of a potentially large number of alerts from heterogeneous sensors. Two important aspects of alert management are alert correlation and alert priority ranking. To the degree that these aspects have been addressed in current systems, heuristic techniques have been used. We describe probabilistic approaches to these critical problems.

The intrusion detection community is actively developing standards for the content of alert messages [3]. Systems adhering to these evolving standards forward attack descriptions and supporting diagnostic data to an alert management interface (AMI), which may be remotely located and consolidating the reports of numerous sensors. In anticipation of these standards, we are evolving systems for alert correlation and ranking. The present development expands on the correlation approach in [3] and provides results of a deployment, and introduces a component that prioritizes alerts. The components rely on statistical similarity measures and probabilistic inference in the form of Bayes networks.

Bayes approaches [4], and probabilistic formalisms in general, represent a minority of inference methodologies employed to date by intrusion detection systems as well as evolving systems for correlating and prioritizing alerts from such systems. Theoretically, a probabilistic system needs to specify the entire joint probability distribution of observable attributes and corresponding priority ranking. This is extremely difficult because of the curse of dimensionality. Instead, the Bayes approach is to assume that dependencies between attributes are local, so a much more compact representation of the system's knowledge base (local conditional probability relations) is possible. The compactness of knowledge representation and the adaptive potential make this approach attractive relative to signature systems.

The remainder of this paper is organized as follows. We first expand on our earlier work in probabilistic alert correlation [3]. We then define the priority ranking mode, including a description of the adaptive training approach, wherein the desired result is to duplicate an expert's priority ranking of a given alert. We then present preliminary results from the alerts generated by various EMERALD monitors [1, 2].

## Sensor Correlation and Alert Fusion

In our view, sensor correlation consists of three key functions, presented in order of increasing difficulty. At the first, we must aggregate large numbers of low-level events (TCP connections, audit records, and so forth) into a manageable number of alert reports. This is achieved by the EMERALD alert thread mechanism. At the next level, it may be desirable for some sensor to be aware of the state of other sensors, and adjust its own state accordingly. Finally, we would like to fuse alerts from heterogeneous sensors, as well as auxiliary information, to provide an operator with intelligently derived meta

alerts. Figure 1 illustrates this hierarchy; the event counts reflect actual counts from a recent 45-day run of some components on our own TCP gateway.

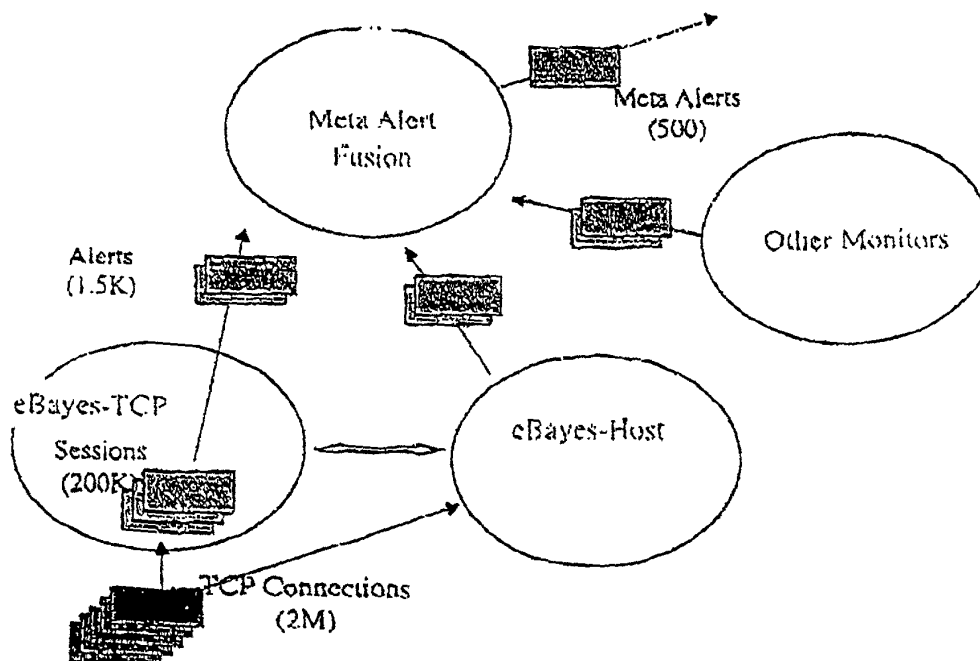


Figure 1: Sensor Correlation Hierarchy

### Alert Threads

EMERAI.D sensors all have the concept of an alert thread, which is compatible with evolving standards for alert content [5]. Sensors may update reports on an existing attack based on transitions to more serious states, significant increases in the attack confidence level, or changes in what the sensor believes the attack to be, among other factors. For extended attacks, these updated reports may flood the analyst's console. By assigning all these reports the same thread ID, a correlation or reporting engine that comprehends threads presents the analyst a single console report, but with the capability to easily drill down if desired.

### Comprehending Multiple Sensors

The system we informally refer to as "eBayes" is actually two sensors: eBayes-TCP, which monitors TCP sessions, and eBayes-Host, which adapts to and monitors valid hosts and services within a protected network. These are coupled in that the former queries the state of the latter, and adjusts accordingly. This adjustment is mathematically straightforward in Bayes formalisms, which start with a prior model of their world and update this model based on the state of numerous features (directly observed or derived) linked to a set of unobservable hypotheses by conditional probability relationships.

The Bayes-Availability monitor adapts to valid hosts and services provided on a protected subnet, and makes available to the session monitor a confidence that a given service requested is valid or not. In the latter case, the session monitor increases its prior suspicion in the case of failed requests. With this coupling, we effectively detect the stealthy portsweeps in the Lincoln Laboratory evaluation data [6]. While the small number of ports in these stealthy attacks is indicative of a specific probe for which a signature can be defined, we taken an approach that we feel is more general and powerful against new probes. The result is that we can detect a variety of other portsweeps (such as nmap and strobe) with no modification of our existing models. Thus, knowledge of the state of one sensor achieves a practical and significant improvement in the sensitivity of another.

The availability monitor also informs our session monitor when a valid service is in a degraded state (again, this is a probabilistic call). Based on this, the monitor alters its prior expectation of the world, adjusting the prior expectation of anomalous values for particular features. In a simulated attack environment, we have a large number of normal clients accessing a network that is under attack. When the attacker achieves a successful denial of service (DOS), these clients suddenly appear anomalous. With our sensor coupling, we detect the DOS and do not raise alerts for the (now failed) legitimate traffic. Essentially, the internal models dynamically come to tolerate values for some features that would normally be indicative of certain attacks. There remains the potential to generate alerts from these sessions, but the evidence must now come from other features. Without this capability, a successful DOS attack might cause hundreds of alerts for the otherwise normal traffic (what we term "collateral damage"). The large number of alerts would overwhelm the analyst, and potentially bury the one valid alert due to the attacker. The approach taken generates one alert for the attack session and one alert that the service is in a degraded state.

We therefore see that our approach, where one sensor comprehends the state of another and adjusts accordingly, both raises sensitivity (evidenced by the ability to detect stealthy portsweeps) and suppresses false (or spurious) alarms. The use of a Bayes formalism makes this coupling mathematically convenient, but is not a requirement.

### Meta Alerts and the Alert Template

We are in the process of defining an enhanced alert template that enriches standards such as IDIP with respect to content while avoiding issues of language specifics. Our template includes the concept of alert thread, which is present in our monitors and alert management components. We include an "anomaly" field in addition to the "confidence" field used in IDIP. We envision these fields being used to condition the reports of some sensors. We also include arrays to describe in greater detail the target(s) of an attack (for example, the specific ports scanned). The API includes functions to allow a downstream component to know which features were set in a particular alert. Finally, we include fields describing the sensor type and placement.

This template is presently used by a diversity of sensors employing both signature and probabilistic techniques. Our early experiments indicate that diverse sensors will be able

to fill this template with content that will be more useful to correlation engines than is currently available.

Our probabilistic alert fusion approach considers feature commonality, feature similarity, and situation-specific expectation of similarity. We maintain a list of "meta alerts" that are possibly composed of several alerts, potentially from heterogeneous sensors. For two alerts (typically a new alert and a meta alert), we begin by identifying features they have in common. Such features include the source of the attack, the target (hosts and ports), the type of the attack, time information, and so forth. With each feature, we have a similarity function that returns a number between 0 and 1, with 1 corresponding to a perfect match. Similarity considers such issues as

- How well do two lists overlap (for example, list of targeted ports)?
- Is one observed value contained in the other (for example, is the target port of a DOS attack one of the ports that was the target of a recent probe)?
- If two source addresses are different, are they likely to be from the same subnet?

Not all sensors produce all possible identifying features. For example, a host sensor provides process ID, while a network sensor does not. Features not common to both alerts are not considered for the overall similarity match.

An important innovation we introduce is expectation of similarity. This is also between 0 and 1, and expresses our prior expectations that the feature should match if the two alerts are related, considering the specifics of each. For example, two probes from the same target might scan the same set of ports on different parts of our subnet (so expectation of matching target IP address is low). Also, some attacks such as SYN FLOOD spoof the source address, so we would allow a match with an earlier probe of the same target even if the source does not match (expectation of match for source IP is low).

We then compose overall alert similarity from feature similarity values normalized by expectation of similarity. The approach is valid regardless of the number of features that overlap, and is thus well suited for use with multiple heterogeneous sensors. The new alert is fused with the meta alert that is most similar, if the similarity is good enough. If no existing meta alerts are a sufficiently good match, the new alert begins a new meta alert thread. Fusion consists of combining features so that the new meta alert is a superset of the previous meta alert and the new observation. Lists are merged, comprehending "hit counts" to each list element. The merge functionality includes "trim functions" to prevent arbitrary list growth. Mathematical details of our similarity functions and expectation of similarity are given in the technical appendix, while the following subsections give more detail on feature fusion and similarity expectation.

The meta alert itself supports the threading concept, so we can visualize composing meta alerts from meta alerts.

## Feature fusion

When the system decides to fuse two alerts, based on aggregate similarity across common features, the fused feature set is a superset of the features of the two alerts. Feature values in fused alerts are typically lists, so alert fusion involves list merging. For

example, suppose a probe of certain ports on some range of the protected network matches in terms of ports with an existing probe that originated from the same attacker subnet, but the target hosts in the prior alert were to a different range of our network. The attacker address list has the new attacker address appended, and the lists of target hosts are merged. The port list matches and is thus unchanged.

Rather than the common practice of listing all feature values ever seen, we maintain as well an associated hit count. For all list features we have a corresponding trim function that trims list elements with extremely low hit counts.

Two important synthetic features are the sensor and thread identifiers of all the component alerts, so that the operator is always able to examine in detail the alerts that contribute to the meta alert report. There is potential for consolidation in a useful sense, but with no loss of information since the component alerts are always available for examination.

### Situation-Specific Similarity Expectation

We now give some examples of how expectation of similarity depends on the situation, that is, the features in the meta alert and the new alert.

If an alert from a sensor has a thread identifier that matches the list of sensor/thread identifiers for some meta alert, the alert is considered a match and fusion is done immediately. In other words, the individual sensor's determination that an alert is an update of or otherwise related to one of its own alerts overrides other considerations of alert similarity.

If the meta alert has received reports from host sensors on different hosts, we do not expect the target host feature to match. If at least one report from a network sensor has contributed to the meta alert and a host sensor alert is received, the expectation of similarity is that the target address of the latter is contained in the target list of the former.

To consider whether an exploit can be plausibly considered the next stage of an attack for which a probe was observed, we expect the target of the exploit (the features host and port) to be contained in the target host and port list of the meta alert.

We declare the attack code similar if we obtain an exact match, but we defer to the attack class in general. For example, in our demonstration environment, we run a variant of mscan that probes certain sensitive ports, that is, it is of the attack class "portsweep". Our host sensors have a specific signature for this attack and call it "mscan". The Bayes sensor trades specificity for generalization capability and has no "mscan" model, but successfully detects this attack as a "portsweep". Given acceptable matches in the target host and port list, these would be considered similar.

If we have observed a probe, we consider a subsequent exploit potentially similar if the target host and service (that is, port) of the exploit are contained in the target host and port lists of the probe. Of course, before fusing the alert, we require similarity with respect to other features as well.

Deciding whether the attacker is similar is somewhat more involved. In the case of an exact match of address, similarity is perfect. We assign high similarity if the subnet

appears to match. In this way, a meta alert may potentially consist of a list of attacker addresses. At this point, we consider similarity based on containment. In addition, if an attacker compromises a host within our network, that host is added to the list of attacker hosts for the meta alert in question. Finally, for attack classes where the attacker's address is likely to be spoofed (for example, the Neptune attack), similarity expectation with respect to attacker address is assigned a low value.

In this fashion, it is possible to recognize a staged attack composed of, for example, a probe followed by an exploit to gain access to an internal machine, and then using that machine to launch an attack against a more critical asset.

## The Priority Ranking Model

Another important aspect of alert management is the need to rank alerts, so that the administrator can concentrate on the most important alerts. We present a Bayes approach to assign a priority ranking to alerts that provide data for several key attributes. Our system has the following features:

- Ability to weight the priority ranking along several attribute groupings, such as attack type or criticality of assets affected
- Compact representation of the influence of the value of an attribute on the priority assigned.
- Incorporation of the administrator's preference profile as to the relative importance of observed values (such as attack type)
- Ranking influenced only by the attributes specified on a given alert; in general, an alert may not report all possible attributes
- Ability to update the prioritization based on observation of a new attribute
- Extensibility of the model to comprehend attributes that may be defined in the future, with minimal perturbation to the rest of the model

Computationally, our approach is to design a Bayes classifier whose output is a priority value and whose observable evidence consists of the attribute values. The influence of an attribute on the output is expressed in terms of conditional probability relations.

The system presented here addresses two problems that arise in IDS alert prioritization. The first is that, while a domain expert can assign a priority to an alert that would generally agree with that assigned by other domain experts, the volume of alerts in a realistic environment with multiple IDS makes it impossible to examine each alert in detail. Second, the depth of intrusion detection domain expertise does not at present exist in most enterprise networks. Our goal is to be able to process a large number of alerts and produce a priority ranking with which a domain expert would agree. The analogy to a "knowledge base" in a Bayes system is the set of conditional probability relations. We have taken the approach of attempting to define these according to our own domain expertise. We recognize that representation of attribute relations as conditional probabilities represents an abstraction from the expert's reasoning process. Therefore, we use this initial representation as a starting point, and provide an adaptive capability

that enables the expert to modify the system's "call". This requires domain knowledge on the part of the expert, but no special knowledge of Bayes systems.

Our model consists of a root node representing the output priority ranking (presently, "low", "medium" and "high"). From the root there are a number of main branches representing attribute groupings. These are linked to the root node by what are effectively "pass through" nodes, whose function is to weight the subtree corresponding to the principal branch. The reader will recall that this differential weighting was noted as one of the principal desirable features of our system. In the present implementation, we have two main branches connecting to nodes representing the influence of the attack attributes on the priority, while the other branch connects to a node representing the influence of asset criticality. Under the first node (node A in Figure 2) are leaf nodes expressing the relationship of attack-specific attributes to the priority. Under the second node (node B in the figure) are leaf nodes representing criticality of asset classes that may be touched by an alert. This representation of attribute groupings as major branches from the root allows for subjective weighting as to the relative importance of the attributes in question. The relative weighting of the major branches is achieved by the "pass through" action of nodes A and B, as discussed in more detail below.

We have proposed this model as a demonstration of concept, namely, that attributes can be grouped and that we wish to potentially express the fact that different groups should have different impacts on the output result. The proposed two attribute groups, described in the next subsections, are only one possible division. We are not limited to two attribute groups, nor are we precluded from a more complex branching structure within the attribute groups. Our proposed model seems to provide a good tradeoff between simplicity and efficacy.

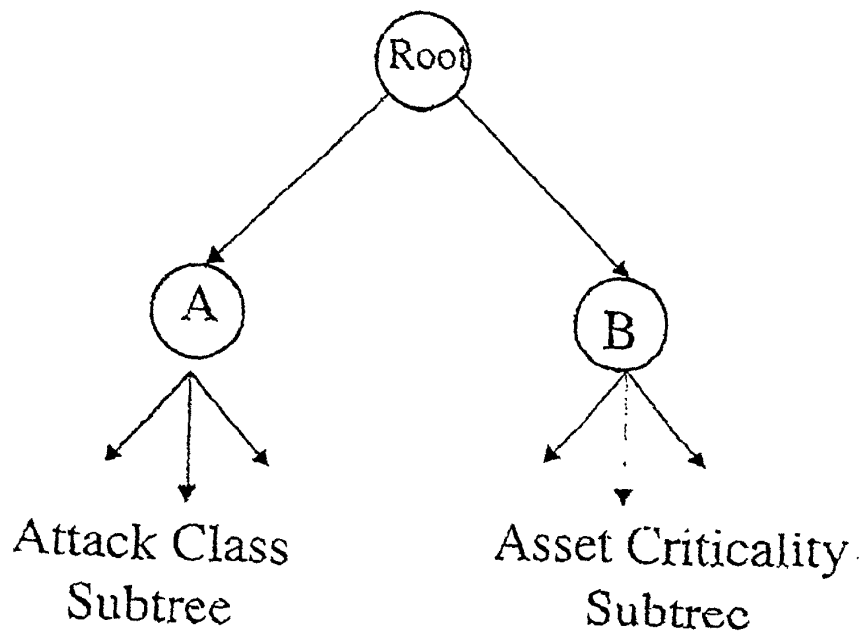


Figure 2: Bayes Model for Alert Prioritization

### Attack Class Attributes

At present, attack class attributes are expressed as a preference profile over the attack classes. This is a configurable measure of concern preassigned to each attack class. We configure our system so that probes are of low concern, while attacks such as unauthorized access are of greater concern. An alert (raw or fused) provides a probability mass over the attack classes; systems capable of a single call assign all their mass to one attack class. This is potentially moderated based both on the sensor's confidence in its call as well as the fusion and prioritization engine's confidence in the sensor. Based on this probability mass, we perform the equivalent of an expected value calculation and return a priority score based on attributes of the attack class. The system has a node that relates this partial score to the overall priority assigned at the root node.

### Asset Criticality Attributes

Figure 3 represents the subtree that relates asset criticality to alert priority. Our initial model identifies five attributes that are potentially critical: user, protocol, service, file (to include directory), and host/subnet. A given alert may not provide values for all these attributes. For example, the TCP session monitor in [2] does not examine user or file. Therefore, for any of these attributes, the alert processor passes to our model one of three values: the attribute was not observed, the attribute was observed and not considered critical, and the attribute was observed and critical. Criticality of an asset is based on a configuration file that reflects security policy. Our model supports dynamic change to the security policy. The elements of the respective CPTs reflect  $P(\text{criticality} = c | \text{priority} = p)$ . Each of these matrices represents two values of criticality by three values of priority. Therefore, the local knowledge base consists of a set of CPTs linking the attribute to the appropriate node on its main branch. If the attribute is not observed in a given alert, the state of the corresponding node is not changed, and thus this attribute does not influence the result one way or the other. If this attribute is observed in a subsequent update for the same alert, our system adjusts the previous prioritization for the new information. Compact representation, use of only the observed attributes, and update of an existing result to comprehend new information are all key desirable features of a system of this type, and they are rigorously addressed in our formalism.






Figure 3: Asset Criticality Subtree

### Pass-Through Nodes

Nodes A and B are the roots of distinct subtrees reflecting the influence of different groups of attributes on the desired result. We can consider these nodes as serving a "pass through" function, propagating the subtree result to the root. If the CPTs relating these major branch nodes to the root are identity matrices, the evaluation from the leaves under the branch is passed through without alteration. Moving mass off the diagonal effectively perturbs and downweights the corresponding subtree result, expressing, for example, different subjective assignments of importance or confidence in a particular attribute group. Initially, these CPTs are near-identity matrices, with some off diagonal mass.

### Adaptive Training

The Bayes inference engine used here has an adaptive capability described in [2]. Briefly stated, the system behaves as if the CPTs are based on effective counts. If a hypothesis (priority assignment) "wins" (posterior belief above a settable learning threshold), entries in the CPT are adjusted slightly in the direction of the observation (which are really the likelihood messages at the leaf nodes). The effective count for the winning hypothesis is aged (multiplied by a decay factor) and incremented by one for the current observation. The effective counts for other hypotheses are aged. Therefore, frequently observed hypotheses approach a saturation count, and as the adjustment depends on the effective count, new observations perturb the current CPT only slightly. Conversely, a very rarely observed hypothesis adapts more quickly to new observations, as its effective count decays to a lower value and thus assigns less weight to past values. This can be thought of as hypothesis-specific annealing.

To utilize this in the present system, we have an interactive facility that randomly generates attribute values and prompts the operator for the priority ranking to be assigned to the alert in question. This becomes the state value for an ancillary "hard call" node, whose value forces the call for the net as a whole. The learning facility is then invoked, and the CPT values adapt accordingly. We believe that with reasonable initial values for the CPT, based on our expert judgment in this area, this iterative adaptation can be achieved with far fewer simulated alerts than would be required to, for example, train a neural net from scratch.

## Results

### Live Data

The following result is from a probe attack detected against our internal network that apparently looked for various ports on different subnets over several days. Each individual component of the attack is detected by the Bayes TCP sensor, which does not know beforehand that this is a sensitive set of ports. The attack does not cause a denial of service.

```

portsweep 1.000 2000-08-10 06:50:28 from 193.230.37.2 ports
1268 to 32434 dt= 0.321
count 164 max age count 0.16 code 3 evc 1 max-err 3.83 -
opn 0 -oip 0 -oport 0
30 dest IPs: aaa.bbb.1.1 aaa.bbb.3.1 aaa.bbb.4.1 aaa.bbb.5.1
aaa.bbb.6.1 aaa.bbb.7.1 aaa.bbb.8.1 aaa.bbb.9.1 aaa.bbb.10.1 aaa\
bbb.11.1 aaa.bbb.12.1 aaa.bbb.13.1 aaa.bbb.14.1 aaa.bbb.15.1
aaa.bbb.16.1 aaa.bbb.17.1 aaa.bbb.18.1 aaa.bbb.19.1 aaa.bbb.20.1 \
aaa.bbb.21.1 aaa.bbb.22.1 aaa.bbb.23.1 aaa.bbb.24.1 aaa.bbb.25.1
aaa.bbb.26.1 aaa.bbb.27.1 aaa.bbb.28.1 aaa.bbb.29.1 aaa.bbb.30.1
aaa.bbb.31.1
6 dest ports: 635(30) 110(29) 143(29) 53(27) 21(27) 109(22)
SEL 0.000 0.000 0.000 0.000 0.000 0.000
0.000 0.000 0.000 0.994 0.004 0.002 0.000
PCODE 0.000 0.000 0.000 0.000 1.000 0.000 0.000
0.000
SVC DIST 0.417 0.000 0.194 0.000 0.369 0.000

```

```

Non-sys alloc ports 6 port_anom 0.990528 code_anom 0.987018
Invalid hosts 30 Invalid ports 6 Neval 24 eval_count 0

```

```

LL-LIST 2000-08-10 06:50:28 aaa.bbb.1.1 to aaa.bbb.31.1 portsweep
1.000

```

Similar attacks with different target IP addresses appeared over several days. The meta alert fusion utility produced the following meta alert:

```

Meta_alert thread 22
Source IPs 193.230.37.2

```

```

Target IPs aaa.bbb.1.1 aaa.bbb.3.1 aaa.bbb.4.1 aaa.bbb.5.1 aaa.bbb.6.1
aaa.bbb.7.1 aaa.bbb.8.1 aaa.bbb.9.1 aaa.bbb.10.1 aaa.bbb.11.1
aaa.bbb.12.1 aaa.bbb.13.1 aaa.bbb.14.1 aaa.bbb.15.1 aaa.bbb.16.1
aaa.bbb.17.1 aaa.bbb.18.1 aaa.bbb.19.1 aaa.bbb.20.1 aaa.\

```

```

bbb.21.1 aaa.bbb.22.1 aaa.bbb.23.1 aaa.bbb.24.1 aaa.bbb.25.1
aaa.bbb.26.1 aaa.bbb.27.1 aaa.bbb.28.1 aaa.bbb.29.1 aaa.bbb.30.1
aaa.bbb.31.1 aaa.bbb.1.2 aaa.bbb.3.2 aaa.bbb.4.2 aaa.bbb.5.2
aaa.bbb.6.2 aaa.bbb.7.2 aaa.bbb.8.2 aaa.bbb.9.2 aaa.bbb.10.2 aaa.bbb.11.2
aaa.bbb.12.2 aaa.bbb.13.2 aaa.bbb.14.2 aaa.bbb.15.2 aaa.bbb.16.2
aaa.bbb.17.2 aaa.bbb.18.2 aaa.bbb.19.2 aaa.bbb.20.2 aaa.bbb.21.2

```

From 2000-08-10 06:50:28 to 2000-08-14 01:47:18

#### Ports

```

Sum Hits 708.999 dot product 0.166775
Index 635 Prob 0.170455
Index 110 Prob 0.168534
Index 143 Prob 0.167341
Index 53 Prob 0.163906
Index 21 Prob 0.169958
Index 109 Prob 0.157806

```

```

Threads :4860 5564 6314 6980 7650 8223 8767 9287 9778 10350 10964 11577
12188 12634 13033 13437 13802 14166 14525 14887 15298 15677 16015 16357
16689 17040 17321 17584 17907 18238 18558 18851 19185 19607 19980

```

#### Attack steps: portswEEP

This sequence of attacks was confirmed as a virtually certain attack by our system administration staff. The list of thread identifiers permits the administrator to examine any of the stages in the attack. In this case, each attack stage is considered a portswEEP; if the stages consisted of different attack classes, these would be listed under "Attack steps" above. This attack was a probe (low priority) but accessed at least one critical host, so the overall assigned priority is 0.45 (on a scale of 0 to 1).

#### Simulation Environment

We describe the operation of the above components in a simulated attack environment. The environment simulates an electronic commerce site, and provides Web and mail services over the Internet. The protected network is behind a firewall, and is instrumented with several host, network, and protocol monitors. Two machines are visible through the firewall, and two auxiliary machines are used for network monitors and alert management. The firewall blocks all services other than Web and mail. We have simulated traffic that appears to come from a number of sources, and an attacker who executes certain well-known attacks.

The attack begins with an mscan probe to the two machines visible through the firewall. Signature-based sensors on the two machines detect this and set the attack code for mscan in their alert report. The Bayes network sensor detects the attack as of the class "portswEEP". The target port lists for these alerts match, as does the attacker address. Similarity with respect to the target address depends on the order in which the alerts arrive (the order at which the alerts arrive is not deterministic as the sensors are distributed). Until the network sensor's alert is received, we do not expect target addresses to match for distinct sensors, since a host sensor can typically report only its own address as a target. We expect the network sensor's target list to contain the list of targets seen so far, and we expect the target address of subsequent host alerts to be

contained in the meta alert's target host list. Regardless of arrival order, these alerts are fused. We have set a preference profile value of LOW (0.1) for probe attacks, but since one critical host and one critical service were probed, the overall alert is prioritized by the Bayes prioritization component at 0.64 on a scale from 0 to 1.

The attacker next uses a CGI exploit to obtain the password file from the Web server. This is detected by the host sensor. The fusion engine considers this a new step in the existing attack because of a match in attacker address as well as the fact that the target host and port are contained in the target and port list of the existing probe. Criticality rises to over 0.95 because of the attack escalation (probe to unauthorized access) and asset criticality (critical host, service, and now file).

The next stage of the attack is a buffer overflow to gain access on the host providing mail service. This is detected by the host sensor. Once again, the attack target is already included in the list of the meta alert, so this alert is fused as well. It is important to note that the address of this compromised host is added to the list of attacker source hosts (that is, it is both a target and now a potential source for an attack).

Although telnet through the firewall is blocked, the attacker may telnet from the compromised internal host to the critical host, and he now does so. On that host, he uses another overflow attack to obtain root access. He is now free to change Web content. All these actions are detected, and fused into one meta alert expressing all stages of the attack. Criticality at this point is extremely high (essentially 1.0).

It is worth pointing out that all the sensors in question provide response directives which, if followed, would stop the above attack in progress.

## Summary

As intrusion detection systems are more widely deployed, the problem of alert management assumes paramount importance. This is because a qualified domain expert is not able to examine all alerts produced on realistic enterprise networks by present IDSs, and at any rate such domain expertise is still not widespread. We have proposed a system that addresses two aspects of alert management, namely, alert correlation and prioritization. The probabilistic approach of these systems gives an alternative and complement to the heuristic techniques more commonly used in intrusion detection.

We have adapted and extended notions from the field of multisensor data fusion for alert correlation. The extensions are principally in the area of generalizing feature similarity functions to comprehend observables in the intrusion detection domain. The approach has the ability to fuse alerts for which the match is good but not perfect. A list of identifiers is maintained with each meta alert that permits linkage to the individual component alerts.

Our approach to prioritization is to construct a Bayes network whose observables are the attributes potentially reported for an alert and whose output is an alert priority. The knowledge base is compactly represented as local conditional probability relations between members of attribute groups and between these attribute groups and the output priority. The system has an adaptive capability wherein a domain expert can modify the system's "call" for randomly generated alert exemplars. This training facility can be

१

This research was sponsored by DARPA under contract number F30602-99-C-1049. The views herein are those of the author(s) and do not necessarily reflect the views of the supporting agency.

[illegible]

## References

1. Porras, P. and Neumann, P. "EMERALD: Event Monitoring Enabling Responses to Anomalous Live Disturbances", National Information Security Conference, 1997.  
<http://www.sdl.sri.com/emerald/emerald-niss97.html>
2. Valdes, A. and Skinner, S. "Adaptive, Model-based Monitoring for Cyber Attack Detection", Recent Advances in Intrusion Detection (RAID 2000), Toulouse, France, October 2000, <http://www.raid-symposium.org/raid2000/program.html>
3. Valdes, A. and Skinner, S. "Blue Sensors, Sensor Correlation, and Alert Fusion", Recent Advances in Intrusion Detection (RAID 2000), Toulouse, France, October 2000, <http://www.raid-symposium.org/raid2000/program.html>
4. Pearl, J. "Probabilistic Reasoning in Intelligent Systems", Morgan-Kaufmann (1988).
5. Erlinger, M. and Stanniford, S. "Intrusion Detection Interchange Format", <http://www.ietf.org/html.charters/idwg-charter.html>
6. Lippmann, Richard P, et al. "Evaluating Intrusion Detection Systems: The 1998 DARPA Off-Line Intrusion Detection Evaluation," Proceedings of DARPA Information Survivability Conference and Exposition, DISCEX'00, Jan 25-27, Hilton Head, SC, 2000, <http://www.ll.mit.edu/IS17/ideval/index.html>

## Technical Appendix: Mathematical Description of Similarity and Similarity Expectation

### Similarity

In the following discussion,  $X$  is defined to be the value of a feature in a new alert, and  $Y$  is the value of that feature in an existing alert class to which it is being compared. For features that have a single value, the features of a new alert and an existing alert class are defined to be similar if their feature values are equal. That is,

$$SIM(X, Y) = \begin{cases} 1.0, & X = Y \\ 0, & \text{otherwise} \end{cases}$$

The more general case is one in which a feature of an alert includes a list of observed values, and a "hit count" of the number of times each value has been observed. For reasons of normalization, the hit count may be converted to a probability. In this sense,  $X$  and  $Y$  are lists of feature and probability values, possibly of different lengths. A probability vector describes a pattern over the observed categories, where

$p_X(C)$  = probability of category  $C$  in list  $X$

$p_Y(C)$  = probability of category  $C$  in list  $Y$

$P_X$  = probability vector over categories observed for  $X$

$P_Y$  = probability vector over categories observed for  $Y$

Note that these probability vectors are more general than the typical binary representation frequently employed in feature matching. The notation  $C \in X$  to denote that category  $C$  occurs in list  $X$ . Then we can use a variety of similarity functions based on normalized dot products (or equivalently, cosine separation of feature vectors) such as

$$Sim(X, Y) = \frac{\left[ \sum_{C \in X \text{ AND } C \in Y} P_X(C) \times P_Y(C) \right]^2}{(P_X \cdot P_X)(P_Y \cdot P_Y)}$$

If the two lists are the same length, then this measure gives the square of the cosine between the two. This has an intuitive geometric property that is not shared by, say, the dot product as is commonly used in the pattern matching literature. If the patterns  $(1, 0)$ ,  $(0, 1)$ , and  $(0.5, 0.5)$  are understood to be over the same two categories, then

$$Sim(\{1, 0\}, \{0, 1\}) = 0$$

$$Sim(\{0.5, 0.5\}, \{0, 1\}) = Sim(\{0.5, 0.5\}, \{1, 0\}) = 0.5$$

In other words, the first two patterns are orthogonal, and the third is halfway in between.

Similarity measures of this type are typically found in multisensor data fusion applications. For our purposes, we must extend this notion of similarity for some features. For example, there are special cases of similarity where we want to consider an observed feature "similar" if it is included in the list of the candidate alert class to which

it is being compared. A probe attack may search for a vulnerable service on some host, and then an attack to a specific host exploits the vulnerability. In this case, we consider the target "similar" if the target of the exploit is included in the target list of the probe. In addition, IP addresses have special similarity functions to express an intermediate value for addresses appearing to originate from the same subnet. Finally, in the case where a target host is compromised, its address is added to the attacker's address list, and is a candidate for similarity evaluation as an attack originator.

### Expectation of Similarity

As was discussed earlier, the nature of an alert may change an expectation of which features of a new alert and an existing alert class should be similar. For example, a syn flood is a type of attack in which the source IP address is typically forged. In this case, similarity in the source IP address would not be considered when assessing the overall similarity between a new alert triggered by a syn flood attack and an existing alert class.

Expectation of similarity is a feature-specific number between 0 and 1, with 1 indicating a strong expectation of a match, and 0 indicating that two otherwise similar alerts are probably not similar with respect to a specific feature. An alert state is represented as a probability vector over candidate states. For each candidate alert state, each feature has a vector of the same length whose elements are the similarity expectation values given the state. This table may be initially populated with MEDIUM (about 0.6) values, except where it is known otherwise. For example, in an access from a compromised host or in the case of a DOS, there is a LOW expectation that the source IP will match that of the attacker.

Based on the alert state, the similarity expectation may be dynamically generated as the weighted sum of the elements of an expectation table, with the weights from the (evolving) alert state distribution. This is itself an expectation in the statistical sense, conditioned on the belief over the present alert state. This is general enough to cover the situation of an unambiguous call from a signature engine, in which case the distribution over states is 1.0 at the state corresponding to the call and 0 elsewhere. Algebraically, for a feature  $J$

$$E_J = \sum_{i \in \{alert\_states\}} BEL(i) E_J(i)$$

$E_J$  = expected similarity for feature  $J$ , given present state

$BEL(i)$  = belief that the attack state is presently  $i$

$E_J(i)$  = element  $i$  of the lookup table of expected similarity for measure  $J$  given state  $i$

Both the new alert and the alert class to which it is being compared each compute the similarity expectation for each feature. Feature similarity and similarity expectation are then combined to form a single value of alert similarity. This is done by combining feature similarities and normalizing by similarity expectation, over the set of common features:



$$SIM(X, Y) = \frac{\sum_j E_j^X E_j^Y SIM(x_j, y_j)}{\sum_j E_j^X E_j^Y}$$

X = new alert

Y = existing alert

J indexes features

$E_j^{X(Y)}$  = similarity expectation for feature J, alert X(Y)

A similar definition can be formed by taking products rather than sums above, and then taking the geometric mean. This has some advantages, but can be overly influenced by a single extremely small value.

Global changes:

' '	to	' '
" "	to	" "
multi-sensor	to	multisensor
port scan	to	portsweep
net	to	network
subnet	to	subnetwork
internet	to	Internet
web	to	Web

Be consistent throughout: portsweep, port sweep, port scan, stealthy portsweep, stealth port sweep,

Be consistent: thread identifier or thread ID

At first use, establish IDS as an abbreviation for intrusion detection system and use thereafter, or spell out in full throughout the paper.

If ID is used for "identifier", do not use for "intrusion detection" (unless as part of IDS).

Be sure that TCP and CPT are not confused, and define CPT at first use.

Check "set of CPT" --should it be "set of CPT's"? Note that set needs a singular verb: "set is"  
Note that CPT needs a singular verb: "the CPT is", but "these CPTs are"

Run spell check.